arXiv:1511.08670v2 [stat.AP] 16 Mar 2016

A Bayesian binary classification approach to pure tone audiometry

Marco Cox
Eindhoven University of Technology
Eindhoven, The Netherlands
m.g.h.cox@tue.nl

Bert de Vries
GN ReSound and Eindhoven University of Technology
Eindhoven, The Netherlands
bdevries@ieee.org

Abstract

The pure tone hearing threshold is usually estimated from responses to stimuli at a set of standard frequencies. This paper describes a probabilistic approach to the estimation problem in which the hearing threshold is modelled as a smooth continuous function of frequency using a Gaussian process. This allows sampling at any frequency and reduces the number of required measurements. The Gaussian process is combined with a probabilistic response model to account for uncertainty in the responses. The resulting full model can be interpreted as a two-dimensional binary classifier for stimuli, and provides uncertainty bands on the estimated threshold curve. The optimal next stimulus is determined based on an information theoretic criterion. This leads to a robust adaptive estimation method that can be applied to fully automate the hearing threshold estimation process.

1 Introduction

The most common and basic way to quantify hearing impairment is to estimate the pure tone hearing threshold (HT), which is called *pure tone audiometry* (PTA) (Yost, 1994). The conventional way to estimate the pure tone HT of a person is to incrementally approximate the lowest tone intensity that can still be perceived at a set of standard frequencies ranging from $250\ \mathrm{Hz}$ to $8\ \mathrm{kHz}$ using a staircase "up $5\ \mathrm{dB}$ - down $10\ \mathrm{dB}$" approach (Carhart and Jerger, 1959). This simple approach however suffers from a number of drawbacks:

- It does not provide an uncertainty measure for the estimated thresholds, making it hard to define an objective stopping criterion.
- It is difficult to exploit potential prior knowledge about the threshold curve that is being estimated.
- It does not use all information available in the data by assuming that the hearing thresholds at different frequencies are uncorrelated. This leads to redundant measurements being required to achieve the desired accuracy.
- Restricting the stimuli to a fixed set of standard frequencies requires additional assumptions to estimate the threshold at other frequencies. The most common solution to this is to assume piecewise linearity of the HT curve on a semi-logarithmic frequency scale, for example in an audiogram. An audiogram depicts the difference between a particular HT curve and the HT curve of a normal-hearing person.

This paper presents a full probabilistic approach to the pure tone hearing threshold estimation problem that addresses these drawbacks. Taking a probabilistic approach allows accounting for uncertainty in the patient's responses as well as incorporating prior knowledge in a fundamental way. Multiple probabilistic methods have already been proposed to address the first two drawbacks, for example in (Barthelmé and Mamassian, 2008). The novelty of our approach is in probabilistically modelling the HT as a continuous function of frequency, rather than using multiple independent response models (psychometric functions) for a discrete set of standard frequencies. This is by endowing the threshold curve with a Gaussian process (GP) prior (Rasmussen and Williams, 2006), which is why we refer to our method as 'GP-PTA'. If the responses are binary (audible or non-audible), the full model can be interpreted as a two-dimensional binary classifier that is specified by a GP in the frequency dimension and by a psychometric (sigmoid) function in the intensity dimension. It provides uncertainty bands on the resulting threshold estimate, which enables fundamental and objective stopping criteria. Moreover, it allows one to find the optimal next stimulus based on the data processed so far, and stimuli at any frequency can be used. These properties reduce the number of trials required to achieve the desired accuracy. Minimizing the amount or required trials is important to reduce the cognitive burden on the (often elderly) patients. Figure 1 depicts the high level structure of the GP-PTA method.

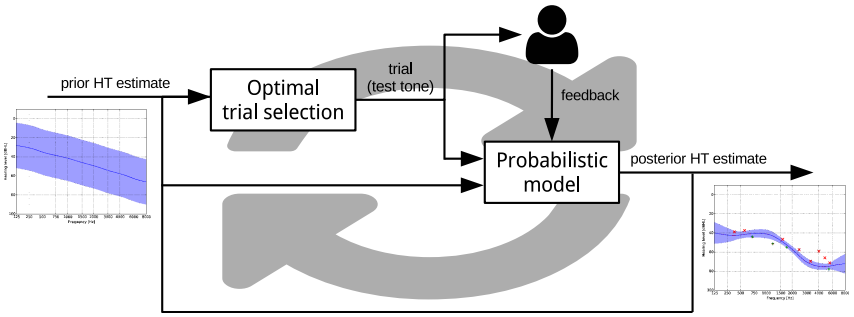

Figure 1: High level overview of the proposed GP-PTA method. The probabilistic model is a two-dimensional binary classifier specified by a GP in the frequency dimension and by a psychometric function in the intensity dimension. The posterior distribution of the threshold curve is inferred incrementally by repeating two steps. First, the most informative next trial (stimulus) is determined based on the data processed so far. Next, Bayesian inference techniques are utilized to find the posterior distribution of the HT curve based on the person's response and the prior distribution.

More specifically, our GP-PTA method combines the following elements:

- A probabilistic response model is specified to capture uncertainty in the patient's responses (Section 2). Accounting for inconsistency in the responses is crucial for providing error bands on the estimated threshold curve.
- A generative probabilistic model is derived by combining the response model with a GP (Section 3). The GP provides a natural way to take prior knowledge about the threshold curve into account, and relaxes the piecewise linear assumption for the threshold curve to a more natural smoothness assumption.
- Bayesian inference techniques are applied to derive an inference algorithm for the probabilistic model. This algorithm is used to 'learn' the threshold curve from the data.
- An active learning loop is obtained by adaptively selecting the optimal next stimulus according to the "Bayesian active learning by disagreement" (BALD) objective function (Section 4).

These elements are combined in simulations to evaluate the performance of the method.

2 Experiment setup and response model

To estimate a patient's pure tone HT curve, an *experiment* consisting of a sequence of *trials* is conducted. A trial consists of the patient labelling a pure tone stimulus with a certain frequency f

and intensity level as either 'audible' or 'non audible'. The sound intensity level is expressed by the *hearing level* h in dB *Hearing Level* (dBHL), which is defined as the intensity relative to the hearing threshold of a person with no hearing deficit (0 dBHL). The patient's response is captured in binary class label y according to:

$$y(f,h) = \begin{cases} +1, & \text{if } (f,h) \text{ is audible,} \\ -1, & \text{otherwise.} \end{cases} \tag{1}$$

Thus, a trial is represented by the 3-tuple (f,h,y). Multiple instances of the same stimulus close to the hearing threshold will not always be labelled consistently due to human inconsistency and contextual changes. To capture this uncertainty in the responses, the 'true' HT is assumed to be evaluated under white Gaussian perceptual noise: $\mathcal{N}(0,\sigma_p^2)$. This assumption leads to a probabilistic response model:

$$\begin{aligned} P(y|f,h) &= \Pr\{y\cdot(h-\mathcal{HT}(f)) > \mathcal{N}(0,\sigma_p^2)\} \\ &= \int_{-\infty}^{y\cdot(h-\mathcal{HT}(f))} \mathcal{N}(h'|0,\sigma_p^2)\mathrm{d}h' \\ &= \Phi\left(\frac{y\cdot(h-\mathcal{HT}(f))}{\sigma_p}\right), \end{aligned} \tag{2}$$

where $\mathcal{HT}$ is the unknown 'true' hearing threshold function and Φ is the cumulative density function of the standard normal distribution. Note that $P(y|f,h)$ is a Bernoulli distribution. This response model assumes that the perceptual noise variance σ_p^2 is independent of frequency, an assumption that can easily be relaxed in the future.

3 PTA by Bayesian binary classification

The response model from (2) can be interpreted as a binary classifier for stimuli (f,h). This binary classifier has a single decision boundary in the frequency-intensity space, defined by the unknown function $\mathcal{HT}(\cdot)$. Let $\mathcal{D}$ denote a data set consisting of N trials:

$$\mathcal{D} \triangleq \{(f,h,y)_1,\ldots,(f,h,y)_N\}.$$

The goal is to use the response model to estimate $\mathcal{HT}$ as accurately as possible from $\mathcal{D}$. A good method will achieve high accuracy with a minimal number of trials. The conventional and most simple way to estimate $\mathcal{HT}$ is to restrict f to a set of standard frequencies, and to estimate $\mathcal{HT}(f)$ at those frequencies by maximizing the likelihood of the data points involving that frequency. However, this approach has serious drawbacks as discussed in the introduction. To address those drawbacks, we propose a probabilistic, Bayesian approach to the estimation problem. This involves endowing the unknown function $\mathcal{HT}$ with a prior to get a full generative model for the data, and then using Bayesian inference techniques to find the posterior distribution over $\mathcal{HT}$. This approach allows processing trials at any frequency, uses all information from the trials, provides uncertainty bands on the final HT estimate, and relaxes the common piecewise linearity assumption for the HT curve to a more realistic smoothness assumption. First, we define the full generative model (Section 3.1) and show how to generate predictions from it. Next, we derive an inference algorithm to find the posterior distribution over $\mathcal{HT}$ (Section 3.2).

3.1 Model definition

To extend the response model from (2) to a full generative model, one has to specify $\mathcal{HT}$ and σ_p. The perceptual noise variance σ_p^2 is assumed to be (approximately) known in this work, but it is important to note that this quantity in principle can also be estimated by Bayesian inference.

3.1.1 Frequency warping

To define a prior on $\mathcal{HT}(f)$, it is convenient to first transform (warp) the frequency domain to a psychoacoustically relevant domain. Defining the prior in a psychoacoustical domain instead of directly in the frequency domain simplifies matters since the smoothness of the HT curve is more constant in the psychoacoustical domain, and our aim is to encode this smoothness property in the

prior on $\mathcal{HT}$. A psychoacoustical domain is defined such that it matches the (nonlinear) human perception of frequency shifts. Multiple psychoacoustical scales are being used in practice, based on different definitions. Popular scales include the semitone scale, Mel scale, and Bark scale (Yost, 1994). The semi-logarithmic horizontal axis of an audiogram is closely related to those scales. In this work, the psychoacoustical domain is defined by the Bark scale, but it is important to note that our method in general is independent of the chosen transformation.

Multiple transformations from frequency to Bark scale have been proposed in the literature, for example in (Traunmüller, 1990, Wang et al., 1991). The transformations differ in optimization criteria, accuracy, and complexity. The GP-PTA model requires an invertible transformation, which is why we use the simple transformation proposed by Wang et al. (1991):

$$\text{bark}(f) \triangleq 6 * \sinh^{-1}\left(\frac{f}{600}\right), \tag{3}$$

with f in Hz. We denote the transformed frequency as $x = \text{bark}(f)$. The unknown 'true' HT function in the transformed domain is denoted as

$$g(x) \triangleq \mathcal{HT}(\text{bark}^{-1}(x)), \tag{4}$$

and so as a consequence $\mathcal{HT}(f) = g(\text{bark}(f))$. In the remainder of this paper, the frequency in the psychoacoustical domain, x, is used instead of f itself.

3.1.2 Full generative model

The prior on $\mathcal{HT}$ is defined indirectly by the prior on g. Unknown function g is endowed with a Gaussian process prior, which defines a distribution over all continuous real-valued functions. A GP is fully specified by a collection of random variables, a mean function, and a covariance (kernel) function. Let

$$g(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k_\theta(\mathbf{x}, \mathbf{x})) \tag{5}$$

denote that function g is drawn from a Gaussian process with mean function $m(\cdot)$ and covariance function $k_\theta(\cdot, \cdot)$. The covariance function is governed by hyperparameters θ. The GP prior is appropriate here because it supports non-linear functions and leads to a tractable (approximate) posterior distribution. The mean and covariance functions encode our prior knowledge about g. In particular, we use the squared exponent covariance function, which encodes a smoothness assumption. For more background on GPs, we refer to (Rasmussen and Williams, 2006).

Since g is a random function, its function values can be represented by random variables. These random variables are denoted by $g_x \triangleq g(x)$ for scalar x, and by $\mathbf{g_x} \triangleq [g(x_1), \ldots, g(x_n)]^T$ for vector $\mathbf{x}$.

The full generative model is now obtained by replacing $\mathcal{HT}$ in the response model (2) by latent function g:

$$P(y, g|x, h) = P(y|g, x, h) \cdot p(g), \tag{6a}$$

$$P(y|g, x, h) = \Phi\left(\frac{y \cdot (h - g(x))}{\sigma_p}\right), \tag{6b}$$

$$p(g) = \mathcal{GP}(m(\cdot), k_\theta(\cdot, \cdot)), \tag{6c}$$

where p denotes a continuous probability distribution and P a discrete one. Note that y is a random variable whereas g is a random function. If σ_p is known, the complete posterior is given by

$$P(y, g|x, h; \mathcal{D}) = P(y|g, x, h) \cdot p(g|\mathcal{D}). \tag{7}$$

Since the first term is independent of $\mathcal{D}$, model inference in this setting amounts to finding the posterior distribution over the latent function, $p(g|\mathcal{D})$. The posterior distribution on $\mathcal{HT}$ itself is then given by $\mathcal{HT}(f) = g(\text{bark}(f))$, which can be evaluated for arbitrary f.

3.1.3 Generating predictions

Once the posterior distribution $p(g|\mathcal{D})$, is known, the generative model can be used to predict the patient's response to any test stimulus (x_*, h_*). The posterior response distribution is given by:

$$P(y_*|x_*, h_*; \mathcal{D}) = \Phi\left(\frac{y_* \cdot (h_* - \mu_*)}{\sqrt{\sigma_p^2 + \sigma_*^2}}\right), \tag{8}$$

where μ_* and σ_*^2 are the mean and variance of the (approximate) posterior GP evaluated at x_*. Appendix A contains the derivation of this result.

3.2 Approximate model inference

A GP prior only yields a GP posterior under Gaussian likelihood (Rasmussen and Williams, 2006). The likelihood $P(y|g,h)$ in the proposed model is not Gaussian, resulting in an intractable posterior $p(g|\mathcal{D})$. A common approach is to approximate the intractable posterior with a Gaussian one using Laplace's method, which we will apply here. Alternative approaches include approximation by expectation propagation, variational optimization, or Markov chain Monte Carlo (MCMC) sampling. The choice for a specific approximate inference algorithm is independent of the model definition.

The Laplace inference algorithm for the GP binary classifier is well known, and is derived for example in (Rasmussen and Williams, 2006). Our model is not a pure GP classifier since the likelihood in (6b) depends on h, which results in a slightly different derivation. Let $\mathbf{x}$, $\mathbf{h}$, and $\mathbf{y}$ be vectors holding the data in $\mathcal{D}$, where $\mathbf{x} = \text{bark}(\mathbf{f})$. The posterior on the latent function, $p(g|\mathcal{D})$, is completely specified by $p(\mathbf{g_x}|\mathcal{D})$ combined with the posterior mean and covariance functions. GP inference consists of two parts: (i) approximating $p(\mathbf{g_x}|\mathcal{D})$ with a Gaussian distribution, and (ii) finding the optimal hyperparameters for the posterior covariance function. The latter is usually done by maximizing the log-marginal likelihood of the data w.r.t. the hyperparameters, which can be applied out-of-the-box to this model. In the remainder of this section we will focus on the non-standard first part.

The true posterior $p(\mathbf{g_x}|\mathbf{x},\mathbf{h},\mathbf{y})$ is approximated by a Gaussian distribution $q(\mathbf{g_x}|\mathbf{x},\mathbf{h},\mathbf{y})$, which is found using the Laplace approximation:

$$q(\mathbf{g_x}|\mathbf{x},\mathbf{h},\mathbf{y}) = \mathcal{N}(\hat{\mathbf{g}}_\mathbf{x}, A^{-1}), \tag{9a}$$

$$\hat{\mathbf{g}}_\mathbf{x} \triangleq \arg\max_{\mathbf{g_x}} p(\mathbf{g_x}|\mathbf{x},\mathbf{h},\mathbf{y}), \tag{9b}$$

$$A \triangleq -\nabla^2_{\mathbf{g_x}} \log p(\mathbf{g_x}|\mathbf{x},\mathbf{h},\mathbf{y})|_{\mathbf{g_x}=\hat{\mathbf{g}}_\mathbf{x}}. \tag{9c}$$

The Laplace approximation reduces the inference problem to the problem of finding $\hat{\mathbf{g}}_\mathbf{x}$ and A, which is worked out in appendix B. Once the approximate posterior has been found, it can be evaluated at any test point x_* using the standard GP formula by integrating out latent function g:

$$q(g_*|\mathbf{x},\mathbf{h},\mathbf{y},x_*) = \int p(g_*|g,x_*)\cdot q(g|\mathbf{x},\mathbf{h},\mathbf{y})\,\mathrm{d}g = \mathcal{N}(\mu_*,\sigma_*^2) \tag{10a}$$

$$\mu_* = m(x_*) + k_\theta(x_*,\mathbf{x})^T k_\theta(\mathbf{x},\mathbf{x})^{-1}\hat{\mathbf{g}}_\mathbf{x}, \tag{10b}$$

$$\sigma_*^2 = k_\theta(x_*,x_*) - k_\theta(x_*,\mathbf{x})^T(k_\theta(\mathbf{x},\mathbf{x}) + W^{-1})^{-1}k_\theta(x_*,\mathbf{x}). \tag{10c}$$

4 Optimal trial selection

In the setting of measuring a patient's HT, one has the ability to pick the next stimulus based on the information gathered so far. An optimal estimation strategy should not only use all available information from the experiment, but also actively select the stimuli whose responses will provide the most information about the quantities of interest. Both contribute to a lower number of trials being required, which reduces the cognitive burden on the (often elderly) person in question. The approach of interleaving optimal trial selection with learning is often referred to as "active learning". This section describes the derivation of a trial selection method based on the "Bayesian Active Learning by Disagreement (BALD)" objective proposed by Houlsby et al. (2011).

Given current data set $\mathcal{D}_n$ and posterior $p(g|\mathcal{D}_n)$, the goal is to find the next stimulus (x_*,h_*) that will provide the most information about latent function g. Once the class label has been observed, the new data point is added to the data set: $\mathcal{D}_{n+1} \leftarrow \mathcal{D}_n \cup (x_*,h_*,y_*)$. Next, the new posterior $p(g|\mathcal{D}_{n+1})$ is approximated using the inference algorithm and the process is repeated until the uncertainty about g drops below the desired threshold. Our goal is to pick the trial that "maximizes the decrease in expected posterior entropy" of g (Houlsby et al., 2011):

$$(x_*,h_*) = \arg\max_{(x,h)} \mathrm{H}[g|x,h,\mathcal{D}_n] - \mathbb{E}_{y\sim P(y|x,h,\mathcal{D}_n)}\mathrm{H}[g|y,x,h,\mathcal{D}_n], \tag{11}$$

where $\mathrm{H}[\cdot|\cdot]$ denotes the conditional Shannon entropy[1]. The objective function from (11) is equivalent to the conditional mutual information[2] of g and y: $\mathrm{I}[g; y|x, h, \mathcal{D}_n]$. Due to the symmetry of mutual information, g and y in (11) can be swapped to get an expression that is easier to compute:

$$(x_*, h_*) = \arg\max_{(x,h)} \mathrm{H}[y|x, h, \mathcal{D}_n] - \mathbb{E}_{g\sim p(g|\mathcal{D}_n)} \mathrm{H}[y|g, x, h]. \tag{12}$$

The rewritten objective only involves conditional entropies of y, which is a binary random variable. Therefore, the general entropy function reduces to the binary entropy function [3]. The first term in (12) is obtained by plugging in the expression for the predictive class distribution under the approximate posterior GP:

$$\mathrm{H}[y|x, h, \mathcal{D}_n] \overset{(17)}{\approx} \mathrm{h}\left(\Phi\left(\frac{h - \mu_x}{\sqrt{\sigma_n^2 + \sigma_x^2}}\right)\right). \tag{13}$$

The second term in (12) is intractable but can be approximated very well by replacing the binary entropy by a squared exponential function as proposed in Houlsby et al. (2011):

$$\begin{aligned}
\mathbb{E}_{g\sim p(g|\mathcal{D}_n)} \mathrm{h}[y|g, x, h] &\overset{(6b)}{\approx} \int \mathrm{h}\left(\Phi\left(\frac{h - g_x}{\sigma_n}\right)\right) \mathcal{N}(g_x|\mu_x, \sigma_x^2) \mathrm{d}g_x \\
&\approx \int \exp\left(-\frac{(h - g_x)^2}{\sigma_n^2 \pi \ln 2}\right) \mathcal{N}(g_x|\mu_x, \sigma_x^2) \mathrm{d}g_x \\
&= \frac{C}{\sqrt{\sigma_x^2 + C^2}} \exp\left(\frac{-(h - \mu_x)^2}{2(\sigma_x^2 + C^2)}\right),
\end{aligned} \tag{14}$$

where $C = \sigma_n \sqrt{\frac{\pi \ln 2}{2}}$. In both terms, the first approximation includes the fact that we approximate the 'true' intractable posterior distribution of g by a GP. Substituting both terms in (12) yields the complete expression for the BALD optimum:

$$(x_*, h_*) = \arg\max_{(x,h)} \mathrm{h}\left(\Phi\left(\frac{h - \mu_x}{\sqrt{\sigma_n^2 + \sigma_x^2}}\right)\right) - \frac{C}{\sqrt{\sigma_x^2 + C^2}} \exp\left(\frac{-(h - \mu_x)^2}{2(\sigma_x^2 + C^2)}\right), \tag{15}$$

where (x, h) should of course be constrained to the domain of interest.

For fixed x, the objective function is maximized by setting $h = \mu_x$ (the decision boundary of the classifier). The proof of this property is contained in appendix C. This result implies that the optimization problem reduces to a one dimensional one:

$$x_* = \arg\min_x \frac{C}{\sqrt{\sigma_x^2 + C^2}} = \arg\max_x \sigma_x^2. \tag{16}$$

So maximizing the BALD objective in this case results in picking the frequency for which the posterior GP variance is largest. However, if the assumption of the response noise being independent of frequency is relaxed, this will no longer be the case. Unfortunately, the objective function is not differentiable w.r.t. x due to the approximate GP inference method that (non-linearly) relates σ_x to x. However, evaluating the objective function is cheap since it only involves an evaluation of the pre-computed approximate posterior GP using (10). Since x is one dimensional, a simple line search works well in practice.

Experiments show that the BALD objective results in frequently performing trials at the borders of the frequency range. Although this is functionally correct, it is not desirable from an audiological point of view. To reflect the fact that the accuracy of the HT estimate is less important at very low and very high frequencies, the objective function may be weighted by a frequency dependent factor.

[1] Shannon entropy is a measure of uncertainty about the value of a random variable: $\mathrm{H}[A] \triangleq \mathbb{E}[-\log p(A)]$. Similarly, a conditional entropy relates to the entropy of a conditional probability distribution: $\mathrm{H}[A|B] \triangleq \mathbb{E}[-\log p(A|B)]$.

[2] Mutual information is a symmetrical measure of dependence between two random variables, which can be expressed in terms of (conditional) entropies: $\mathrm{I}[A; B] = \mathrm{I}[B; A] = \mathrm{H}[A] - \mathrm{H}[A|B]$.

[3] The binary entropy function $\mathrm{h}(\cdot)$ relates the Bernoulli parameter of a binary random variable to its Shannon entropy: $A \sim \text{Bernoulli}(p) \Rightarrow \mathrm{H}[A] = \mathrm{h}(p) \triangleq -p \log(p) - (1 - p) \log(1 - p)$.

5 Related work

Existing pure tone hearing threshold estimation methods range from the simple empirical staircase scheme (Carhart and Jerger, 1959) to more theory-based methods such as QUEST (Watson and Pelli, 1983) and CAST (Özdamar et al., 1990). An adaptive Bayesian method for estimating the parameters of a flexible psychometric function is described in (Barthelmé and Mamassian, 2008). In (de Vries et al., 2010), a Gaussian mixture model is used to model the HT curve and to take prior knowledge into account. The main difference between these existing methods and our GP-PTA method is that the existing ones all focus on estimating the parameters of some psychometric function at a fixed set of frequencies, and ignore the correlation between these parameters at different frequencies. In the context of our framework, this would correspond to estimating the parameters of multiple independent response models: one for every frequency of interest. A recent poster titled "Optimizing Pure-Tone Audiometry Using Machine Learning" (Barbour, 2015) describes a method similar to GP-PTA, in which the continuous HT curve is also modelled by a Gaussian process. Due to the limited amount of details, we are unable to determine the exact amount of similarity between their work and ours.

6 Simulations

The response model, inference method, and trial selection method are combined to incrementally estimate an unknown hearing threshold curve. The approach is tested in simulations using fictional but representative HT curves. To obtain a reasonable prior for the threshold curve, we use a set of standard audiograms (Bisgaard et al., 2010). More specifically, a linear prior with fixed variance is obtained by performing linear regression in the warped frequency domain on the average audiogram resulting from the set of standard audiograms. The simple squared exponential kernel is used in the Gaussian process prior to encode the smoothness assumption on the threshold curve. The kernel hyperparameters are optimized after every trial by maximizing the approximate log marginal likelihood w.r.t. the hyperparameters.

Figure 2 depicts the incremental estimation of a HT curve under response noise standard deviation of 2 dBHL. As is to be expected, increasing the response noise variance leads to slower convergence of the estimate due to more 'false' class labels. Figure 3 depicts the evolution of the mean value of the BALD objective function over time. Ideally, this curve should be monotonically decreasing since every trial contains at least some information. However, the fact that the hyperparameters are optimized using maximum likelihood estimation can cause the mean BALD score to increase on some occasions.

7 Discussion

Taking a probabilistic approach to the problem of estimating the pure tone hearing threshold curve has the important advantage that all forms of uncertainty can be accounted for in a fundamental and objective way. This removes the need for empirically optimized but rather 'ad-hoc' approaches like "up 5 dB - down 10 dB". Moreover, it provides an optimal strategy for selecting the next trial. By accounting for the uncertainty in the patient's responses and in the prior knowledge about the threshold curve, uncertainty bands on the estimated threshold curve can be obtained in a fundamental way. Taking uncertainty in the responses into account is also important from a psychological point of view. Confidence in the correctness of the final estimate should increase if the user can rely on the estimation method to handle answers like "I think I heard it, but I'm not completely sure" are handled in a robust and objective way. The GP-PTA method described in this paper assumes binary answers, and all uncertainty about the response is captured in the response model. However, it is easy to extend the response model to allow more subtle answers, for example by introducing a neutral option. Real-life experiments are required to investigate the usefulness of such extensions. The sigmoid function in our response model takes the role of psychometric function, and follows directly from the the perceptual noise assumptions. Instead of fine-tuning the parameters of the psychometric function itself, as is described for example in (Barthelmé and Mamassian, 2008), we capture the uncertainty about the psychometric function in the prior on the threshold curve.

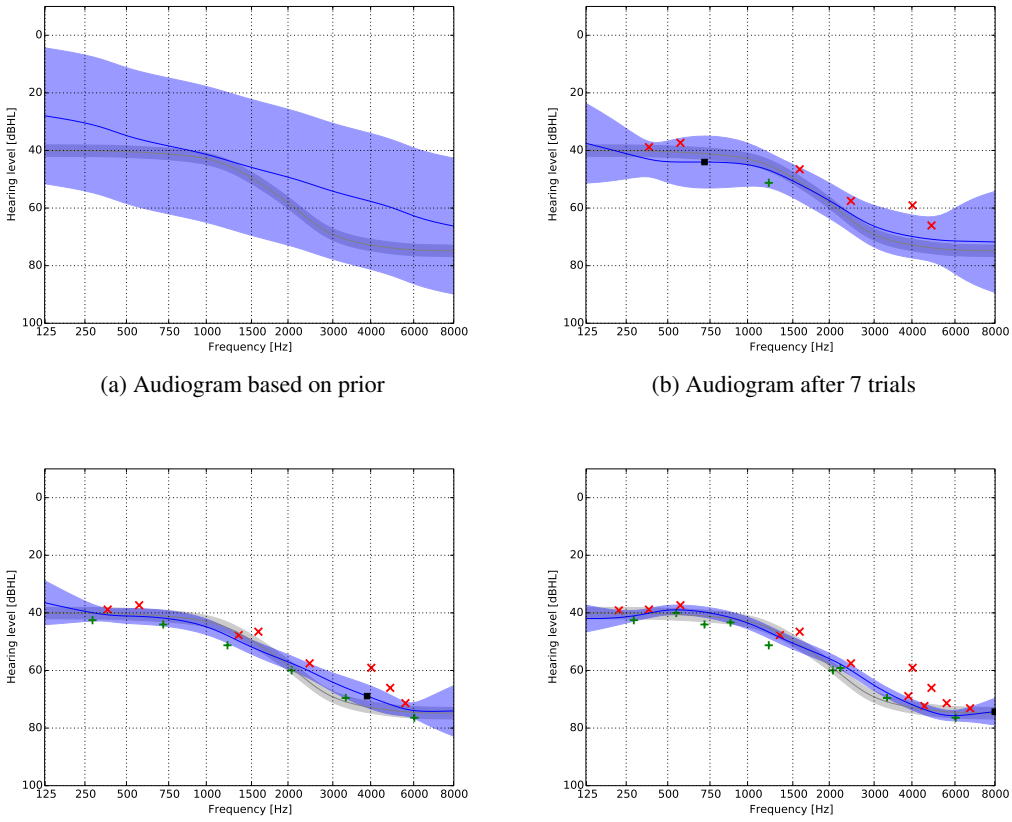

(a) Audiogram based on prior

(b) Audiogram after 7 trials

(c) Audiogram after 14 trials

(d) Audiogram after 21 trials

Figure 2: Incremental estimation of a fictional but representative hearing threshold curve with fixed response noise variance. The shaded areas represent a single standard deviation. The pluses and crosses indicate audible and non audible stimuli, respectively. The square indicates the proposed next trial.

We show that a GP prior on the threshold curve is practical, since it provides a distribution over all possible continuous threshold curves. Moreover, it combines a realistic smoothness assumption with a posterior that can can be approximated analytically. The assumption that the prior perceptual noise variance is known is not unrealistic, since it can be derived from (large) data sets like we did for the simulations. The assumption that this variance is independent of frequency might not be as realistic. However, it is straightforward to extend the model to support frequency-dependent perceptual noise variance. In that case, the BALD criterion does not simplify to picking trials at the frequency with the highest uncertainty about the threshold estimate, which is why we include the complete derivation.

The proposed GP-PTA method provides automatic procedures for selecting the best next trial as well as updating the estimated pure-tone hearing threshold on the basis of patient trial response. As a result, our method in principle supports automated fitting of dynamic range compressing circuits in hearing aids without need for professional human assistance. Moreover, the uncertainty bands on the fitted parameters might be exploited by in-situ tuning algorithms for the hearing aid. Automating these processes is important to be able to cost-effectively cope with the number of hearing impaired persons that is growing worldwide at an alarming rate (Kochkin, 2005, Lin et al., 2011). According to the world health organization, the majority of the 360 million people with disabling hearing loss live in low- and middle-income countries, and the current production of hearing aids meets less than 10% of the global need (World Health Organization, 2015).

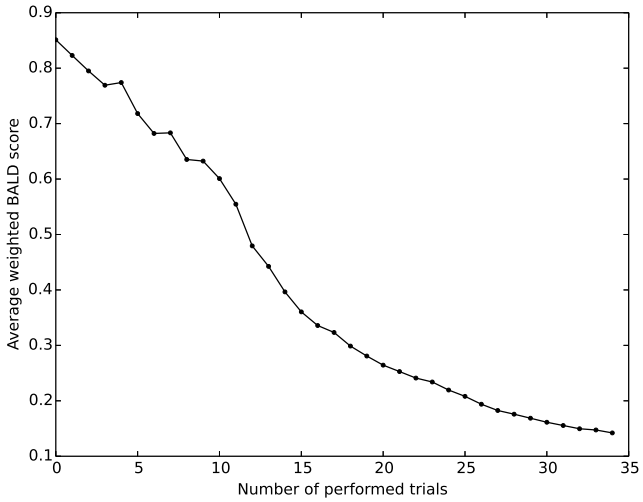

Figure 3: Evolution of the mean of the trial selection objective function (BALD).

Acknowledgements

We like to thank Thijs van de Laar and Tjalling Tjalkens for the interesting discussions, and the GN ReSound research department for scientific support.

Appendix A Generating predictions

The complete posterior (7) of the generative model is converted into a binary classifier $(x, h) \rightarrow y$ by marginalizing out latent function g. If the posterior Gaussian process $p(g|\mathcal{D})$ is known, the posterior response distribution for a test stimulus (x_*, h_*) is given by:

$$\begin{aligned} P(y_*|x_*, h_*; \mathcal{D}) = & \int P(y_*, g|x_*, h_*; \mathcal{D}) \, \mathrm{d}g \\ \overset{(6)}{=} & \int \Phi\left(\frac{y_* \cdot (h_* - g(x_*))}{\sigma_p}\right) \cdot p(g|\mathcal{D}) \, \mathrm{d}g \\ = & \int \Phi\left(\frac{y_* \cdot (h_* - g_*)}{\sigma_p}\right) \cdot \left[\int p(g_*|g, x_*) \cdot p(g|\mathcal{D}) \, \mathrm{d}g\right] \, \mathrm{d}g_* \\ \overset{\text{Laplace}}{\approx} & \int \Phi\left(\frac{y_* \cdot (h_* - g_*)}{\sigma_p}\right) \cdot \mathcal{N}(g_*|\mu_*, \sigma_*^2) \, \mathrm{d}g_* \\ = & \Phi\left(\frac{y_* \cdot (h_* - \mu_*)}{\sqrt{\sigma_p^2 + \sigma_*^2}}\right), \end{aligned} \tag{17}$$

where g_* denotes the latent function value $g(x_*)$. In general, the above integral is not analytically tractable since $p(g|\mathcal{D})$ has no exact analytical solution due to the non-Gaussian likelihood. In this paper, the intractable (non-GP) posterior is approximated by a GP using the Laplace approximation. Under this approximation, the posterior distribution of g_* is Gaussian, and the integral can be solved analytically (see Section 3.9 of Rasmussen and Williams (2006) for the detailed derivation).

Appendix B Laplace approximation of the posterior GP

Laplace's method can be applied to find a Gaussian approximation to an intractable distribution. In this case, the intractable posterior $p(\mathbf{g_x}|\mathbf{x}, \mathbf{h}, \mathbf{y})$ is approximated by a multivariate Gaussian distribution $q(\mathbf{g_x}|\mathbf{x}, \mathbf{h}, \mathbf{y})$ according to:

$$q(\mathbf{g_x}|\mathbf{x}, \mathbf{h}, \mathbf{y}) = \mathcal{N}(\hat{\mathbf{g}}_\mathbf{x}, A^{-1}), \tag{18a}$$

$$\hat{\mathbf{g}}_\mathbf{x} \triangleq \arg\max_{\mathbf{g_x}} p(\mathbf{g_x}|\mathbf{x}, \mathbf{h}, \mathbf{y}), \tag{18b}$$

$$A \triangleq -\nabla^2_{\mathbf{g_x}} \log p(\mathbf{g_x}|\mathbf{x}, \mathbf{h}, \mathbf{y})|_{\mathbf{g_x}=\hat{\mathbf{g}}_\mathbf{x}}. \tag{18c}$$

We will derive an algorithm for finding $\hat{\mathbf{g}}_\mathbf{x}$ and A. Applying Bayes' rule yields:

$$p(\mathbf{g_x}|\mathbf{x}, \mathbf{h}, \mathbf{y}) = \frac{p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) p(\mathbf{g_x}|\mathbf{x})}{p(\mathbf{y}|\mathbf{x}, \mathbf{h})}, \tag{19}$$

where the denominator is independent of $\mathbf{g_x}$. Since prior $p(g)$ is a GP, $p(\mathbf{g_x}|\mathbf{x})$ is Gaussian. Without loss of generality, a zero mean function is assumed for simplicity (the mean function can always be absorbed in the likelihood function):

$$p(\mathbf{g_x}|\mathbf{x}) = \mathcal{N}(\mathbf{0}, k_\theta(\mathbf{x}, \mathbf{x})) = \mathcal{N}(\mathbf{0}, K_{\mathbf{xx}}). \tag{20}$$

Taking the logarithm of the numerator of (19) and writing out the logarithm of the Gaussian prior results in:

$$\begin{aligned} \Psi(\mathbf{g_x}) & \triangleq \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) + \log p(\mathbf{g_x}|\mathbf{x}) \\ & = \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) - \tfrac{1}{2} \mathbf{g_x}^T K_{\mathbf{xx}}^{-1} \mathbf{g_x} - \tfrac{1}{2} \log |K_{\mathbf{xx}}| - \tfrac{N}{2} \log 2\pi. \end{aligned} \tag{21}$$

To find $\hat{\mathbf{g}}_\mathbf{x}$ and A, the first and second order derivatives of Ψ w.r.t. $\mathbf{g_x}$ are required:

$$\nabla_{\mathbf{g_x}} \Psi(\mathbf{g_x}) = \nabla_{\mathbf{g_x}} \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) - K_{\mathbf{xx}}^{-1} \mathbf{g_x}, \tag{22a}$$

$$\nabla^2_{\mathbf{g_x}} \Psi(\mathbf{g_x}) = \nabla^2_{\mathbf{g_x}} \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) - K_{\mathbf{xx}}^{-1} = -W - K_{\mathbf{xx}}^{-1}, \tag{22b}$$

where W is a diagonal matrix since the likelihood factors over the data points. The log-likelihood is defined by the response model from (6b):

$$\log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) = \log \Phi \left(\frac{(\mathbf{h} - \mathbf{g_x})\mathbf{y}}{\sigma_p} \right). \tag{23}$$

The first and second order derivatives of the log likelihood w.r.t. $\mathbf{g_x}$ are required in (22), and we state them here without derivation:

$$\nabla_{\mathbf{g_x}} \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) = \frac{\mathbf{y}\mathcal{N}(\mathbf{g_x}|\mathbf{h}, \Sigma_p)}{\Phi\left(\frac{(\mathbf{g_x}-\mathbf{h})\mathbf{y}}{\sigma_p}\right)}, \tag{24a}$$

$$\nabla^2_{\mathbf{g_x}} \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) = \frac{-\mathcal{N}(\mathbf{g_x}|\mathbf{h}, \Sigma_n)^2}{\Phi\left(\frac{(\mathbf{g_x}-\mathbf{h})\mathbf{y}}{\sigma_n}\right)^2} - \frac{\mathbf{y}(\mathbf{g_x} - \mathbf{h})\mathcal{N}(\mathbf{g_x}|\mathbf{h}, \Sigma_n)}{\Phi\left(\frac{(\mathbf{g_x}-\mathbf{h})\mathbf{y}}{\sigma_n}\right)}, \tag{24b}$$

where $\Sigma_n = \sigma_n^2 I_N$, a $N \times N$ diagonal matrix. Finally, $\hat{\mathbf{g}}_\mathbf{x}$ is found by setting the first order derivative of Ψ to zero:

$$\nabla_{\mathbf{g_x}} \Psi(\hat{\mathbf{g}}_\mathbf{x}) = \mathbf{0} \Longrightarrow \hat{\mathbf{g}}_\mathbf{x} = K_{\mathbf{xx}} \cdot \nabla_{\mathbf{g_x}} \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}). \tag{25}$$

This non-linear equation can be solved by a Newton-Rhapson algorithm. Analog to the derivation in Section 3.4.1 of (Rasmussen and Williams, 2006), the Newton steps are given by:

$$\begin{aligned} \mathbf{g_x}^{\text{new}} &= \mathbf{g_x} - (\nabla^2_{\mathbf{g_x}} \Psi)^{-1} \cdot \nabla_{\mathbf{g_x}} \Psi \\ &= \mathbf{g_x} + (K^{-1}_{\mathbf{x},\mathbf{x}} + W)^{-1} \cdot [\nabla_{\mathbf{g_x}} \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x}) - K^{-1}_{\mathbf{x},\mathbf{x}} \mathbf{g_x}] \\ &= (K^{-1}_{\mathbf{x},\mathbf{x}} + W)^{-1} [W\mathbf{g_x} + \nabla_{\mathbf{g_x}} \log p(\mathbf{y}|\mathbf{h}, \mathbf{g_x})]. \end{aligned} \tag{26}$$

Repeating (26) until convergence of $\mathbf{g_x}$ yields $\hat{\mathbf{g}}_\mathbf{x}$. The total Laplace approximation is obtained by substituting the results in the definition of (18):

$$p(\mathbf{g_x}|\mathbf{x}, \mathbf{h}, \mathbf{y}) \approx q(\mathbf{g_x}|\mathbf{x}, \mathbf{h}, \mathbf{y}) = \mathcal{N}(\hat{\mathbf{g}}_\mathbf{x}, (K^{-1}_{\mathbf{x},\mathbf{x}} + W)^{-1}). \tag{27}$$

Appendix C Proof of BALD optimum at $h = \mu_x$

Theorem. *The BALD objective function*

$$\mathrm{BALD}(x, h) = \mathrm{h}\left(\Phi\left(\frac{h - \mu_x}{\sqrt{\sigma_n^2 + \sigma_x^2}}\right)\right) - \frac{C}{\sqrt{\sigma_x^2 + C^2}} \exp\left(\frac{-(h - \mu_x)^2}{2(\sigma_x^2 + C^2)}\right)$$

is maximized w.r.t. h by $h = \mu_x$. Constant C is given by $C = \sigma_n \sqrt{\frac{\pi \ln 2}{2}}$, $\sigma_n \geq 0$, and $\sigma_x \geq 0$.

Proof. We approximate the binary entropy term up to $\mathcal{O}(h^4)$ by a squared exponential function as proposed in Houlsby et al. (2011):

$$\mathrm{h}\left(\Phi\left(\frac{h - \mu_x}{\sqrt{\sigma_n^2 + \sigma_x^2}}\right)\right) \approx \exp\left(\frac{-(h - \mu_x)^2}{\pi \ln 2(\sigma_n^2 + \sigma_x^2)}\right).$$

Substituting this approximation yields the approximate (smooth and differentiable) objective function Ψ:

$$\mathrm{BALD}(x, h) \approx \Psi(x, h) \triangleq \exp\left(\frac{-(h - \mu_x)^2}{\pi \ln 2(\sigma_n^2 + \sigma_x^2)}\right) - \frac{C}{\sqrt{\sigma_x^2 + C^2}} \exp\left(\frac{-(h - \mu_x)^2}{2(\sigma_x^2 + C^2)}\right).$$

To find the maxima, we differentiate Ψ w.r.t. to h:

$$\begin{aligned} \frac{\partial \Psi}{\partial h} =& \frac{-2}{\pi \ln 2(\sigma_n^2 + \sigma_x^2)} \cdot (h - \mu_x) \cdot \exp\left(\frac{-(h - \mu_x)^2}{\pi \ln 2(\sigma_n^2 + \sigma_x^2)}\right) \\ &+ \frac{2\sigma_n \sqrt{\pi \ln 2}}{(\pi \ln 2 \sigma_n^2 + 2\sigma_x^2)^{\frac{3}{2}}} \cdot (h - \mu_x) \cdot \exp\left(\frac{-(h - \mu_x)^2}{\pi \ln 2 \sigma_n^2 + 2\sigma_x^2}\right), \end{aligned}$$

where constant C has been written out. It is easily verified that $\frac{\partial \Psi}{\partial h}|_{h=\mu_x} = 0$ and that Ψ has a local maximum there. We are left with the problem of proving that this is the global maximum. Another extreme point $h_* \neq \mu_x$ will have to satisfy $\frac{\partial \Psi}{\partial h}|_{h=h_*} = 0$.

$$\left.\frac{\partial \Psi}{\partial h}\right|_{h=h_* \in \mathbb{R} \setminus \mu_x} = 0 \Rightarrow \frac{1}{(h_* - \mu_x)} \left.\frac{\partial \Psi}{\partial h}\right|_{h=h_*} = 0.$$

Taking the logarithm and simplifying yields:

$$(h_* - \mu_x)^2 \underbrace{\left[\frac{1}{\pi \ln 2(\sigma_n^2 + \sigma_x^2)} - \frac{1}{\pi \ln 2 \sigma_n^2 + 2\sigma_x^2}\right]}_{P \leq 0} = \underbrace{\ln \left[\frac{2(\pi \ln 2 \sigma_n^2 + 2\sigma_x^2)^{\frac{3}{2}}}{\pi \ln 2(\sigma_n^2 + \sigma_x^2) 2\sigma_n \sqrt{\pi \ln 2}}\right]}_{Q},$$

which is a second order polynomial in $(h_* - \mu_x)$. Since coefficient P is negative, the equality can only be satisfied if $Q < 0$, a necessary condition for the existence of another extreme point of Ψ. Simplifying Q yields:

$$\begin{aligned} Q =& \ln(2) + \frac{3}{2}\ln(\pi \ln 2 \sigma_n^2 + 2\sigma_x^2) - \frac{3}{2}\ln(\pi \ln 2) - \ln(2) - \ln(\sigma_n) - \ln(\sigma_n^2 + \sigma_x^2) \\ =& \frac{3}{2}\ln\left(1 + \frac{2}{\pi \ln 2}\frac{\sigma_x^2}{\sigma_n^2}\right) - \ln\left(1 + \frac{\sigma_x^2}{\sigma_n^2}\right) \\ =& \frac{3}{2}\ln\left(1 + \frac{2}{\pi \ln 2} u\right) - \ln(1+u), \end{aligned}$$

where we used the substitution $u \triangleq \frac{\sigma_x^2}{\sigma_n^2} \geq 0$. It is immediately clear that $Q = 0$ in case $u = 0$. For the case $u > 0$ we turn to the first order derivative of Q w.r.t. u:

$$\begin{aligned} \frac{\partial Q}{\partial u} =& \frac{3}{2} \cdot \frac{1}{1 + \frac{2}{\pi \ln 2} u} \cdot \frac{2}{\pi \ln 2} - \frac{1}{1+u} \\ =& \frac{u + 3 - \pi \ln 2}{2u^2 + (\pi \ln 2 + 2)u + \pi \ln 2} > 0, \end{aligned}$$

where we used the observation that both the numerator and denominator are positive for $u \geq 0$. Combined with the fact that $Q = 0$ in case $u = 0$, it follows that $Q \geq 0$, which violates the necessary condition for the existence of another extreme point of Ψ. The consequence is that Ψ has just one maximum for fixed x, at $h = \mu_x$. □

References

Barbour, D. (**2015**). "Optimizing Pure-Tone Audiometry Using Gaussian Processes," Proc. 38th Annual MidWinter Meeting, Association for Research in Otolarynology, p. 38.

Barthelmé, S. and Mamassian, P. (**2008**). "A flexible Bayesian method for adaptive measurement in psychophysics," `arXiv:0809.0387 [stat]`.

Bisgaard, N., Vlaming, M. S. M. G., and Dahlquist, M. (**2010**). "Standard audiograms for the IEC 60118-15 measurement procedure," Trends in Amplification 14, 113–120.

Carhart, R. and Jerger, J. (**1959**). "Preferred method for clinical determination of pure-tone thresholds," J. of Speech & Hearing Disorders 24, 330–345.

de Vries, A., Stadler, S., Leijon, A., Dijkstra, T., and Ypma, A. (**2010**). "Efficient evaluation of hearing ability," US Patent App. 12/429,783 (7 October 2010).

Houlsby, N., Huszár, F., Ghahramani, Z., and Lengyel, M. (**2011**). "Bayesian active learning for classification and preference learning," `arXiv:1112.5745 [cs, stat]`.

Kochkin, S. (**2005**). "MarkeTrak VII: Hearing loss population tops 31 million people," The Hearing Review 12(7), 16–29.

Lin, F. R., Niparko, J. K., and Ferrucci, L. (**2011**). "Hearing loss prevalence in the United States," Archives of Internal Medicine 171(20), 1851–1853.

Özdamar, Ö., Eilers, R. E., Miskiel, E., and Widen, J. (**1990**). "Classification of audiograms by sequential testing using a dynamic Bayesian procedure," J. Acoust. Soc. Am. 88(5), 2171–2179.

Rasmussen, C. E. and Williams, C. K. I. (**2006**). *Gaussian processes for machine learning* (MIT Press, Cambridge), 248 pp.

Traunmüller, H. (**1990**). "Analytical expressions for the tonotopic sensory scale," J. Acoust. Soc. Am. 88(1), 97–100.

Wang, S., Sekey, A., and Gersho, A. (**1991**). "Auditory distortion measure for speech coding," Proc. 1991 Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP-91), 493–496.

Watson, A. B. and Pelli, D. G. (**1983**). "Quest: A Bayesian adaptive psychometric method," Perception & psychophysics 33(2), 113–120.

World Health Organization (**2015**). *Fact sheet. Deafness and hearing loss.* Available at *http://www.who.int/mediacentre/factsheets/fs300/en/* (accessed on July 17, 2015).

Yost, W. A. (**1994**). *Fundamentals of hearing: An introduction (3rd ed.), Vol. XIII* (Academic Press, San Diego), 326 pp.